Optimal channel choice for collaborative ad-hoc dissemination

Hu, Liang; Boudec, J-Y. L.; Vojnovic, M.

Published in:
proceedings INFOCOM

Link to article, DOI:
[10.1109/INFOCOM.2010.5462163](https://doi.org/10.1109/INFOCOM.2010.5462163)

Publication date:
2010

Document Version
Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

Citation (APA):
Hu, L., Boudec, J-Y. L., & Vojnovic, M. (2010). Optimal channel choice for collaborative ad-hoc dissemination. In *proceedings INFOCOM* (pp. 1-9). IEEE. <https://doi.org/10.1109/INFOCOM.2010.5462163>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Optimal Channel Choice for Collaborative Ad-Hoc Dissemination

Liang Hu*, Jean-Yves Le Boudec[†], Milan Vojnović[‡]
* Technical University of Denmark [†] EPFL [‡] Microsoft Research

Abstract—Collaborative ad-hoc dissemination has been proposed as an efficient means to disseminate information among devices in wireless ad-hoc networks. It is based on each device forwarding channels that the user of this device is subscribed to and helping forward some other channels. We consider the case where devices have limited resources and thus have to decide which channels to help. The goal is to identify a channel selection strategy that optimizes a global system welfare that is a function of the dissemination times across all distinct channels. We consider a random mixing mobility model under which the channel dissemination time is a function of the number of nodes that forward this channel. We show that maximizing system welfare is equivalent to an assignment problem whose solution can be obtained by a centralized greedy algorithm. We provide empirical evidence that the difference between the system welfare of an optimum assignment and some heuristics proposed in the past can be significant. Furthermore, we show that the optimum social welfare can be approximated by a decentralized algorithm based on Metropolis-Hastings sampling and give a variant that also accounts for the battery energy. Our work provides guidelines how to design decentralized channel selection algorithms that optimize an a priori defined global objective.

I. INTRODUCTION

Several applications that rely on opportunistic data transfers between devices¹ have been proposed recently. In [1], the authors propose a wireless ad-hoc podcasting system, where in addition to downloading content onto devices while connected to a desktop computer, the content is also exchanged between devices while users are on the go. The authors propose several heuristics for prioritization of content exchange between pairs of devices that is based on inferred preference of the user owning the device and that of encountered devices. Another example is CarTorrent [2], a BitTorrent-style content dissemination system that deploys some content solicitation strategies designed to exploit the broadcast nature of wireless medium.

We refer to a stream of content that is generated over time as a *channel*. For example, this could be a podcast feed as well as a profile page of an online social network application such as Facebook or Twitter. While many such services can be provisioned at mobile devices by access to a cloud service through the infrastructure (e.g. WiFi or 3G wireless), it is still of interest to enhance the information dissemination via device encounters. An efficient information dissemination leveraging both the access to the infrastructure and information exchange between devices would well support various mobile content sharing applications, e.g. Serendipity [3], in particular, in environments where access to the infrastructure is intermittent

because of a limited connectivity or costly access to the Internet (e.g. roaming).

We are interested in scenarios where users are willing to devote a limited amount of resources to help disseminate the channels. We assume that each device decides to help a set of channels in addition to channels subscribed by the user of this device. In some scenarios, the number of channels forwarded per user would be a small portion of the total number of distinct channels. For example, there are several thousands of distinct podcast feeds accessible through the Zune Social online service [4] while typically only a few subscribed channels per user [5]. Having a limit on the number of channels helped by a device limits the cost per device such as bandwidth usage, energy consumption, and storage. We consider cooperative scenarios where each device implements a channel selection strategy that optimizes a global system welfare objective; such cooperative setting also underlies previous work [1]. Cooperation could be enforced by the protocol design or induced by use of some incentive schemes.

We would like to formulate a channel selection strategy that optimizes a system welfare objective that is a function of channel dissemination times. The key assumption that facilitates our formulation is that the dissemination time of any channel is a function of the fraction of devices that forward this channel. For example, this assumption holds for the random mixing mobility model, where every device encounters any other device equally likely. For more general mobility, the relation between the dissemination time and the fraction of forwarding devices could be characterized by inference. We would like to emphasize that in this paper we do not advocate any specific function for the relation between the dissemination time and the fraction of forwarding devices; a thorough analysis of this is left for future work. We formulate a system welfare problem where the objective is to optimize a sum of the utility functions of dissemination times of individual channels. We show that for a broad class of utility functions, optimizing the system welfare is equivalent to an assignment problem whose solution can be obtained by a centralized greedy algorithm. We provide empirical evidence based on subscription profiles to podcast feeds of Zune Social [4], [5] that there is a substantial difference between the optimum social welfare and that of some previously-proposed heuristics.

Furthermore, we consider the problem of defining a practical and decentralized algorithm run by individual devices to optimize the system welfare. We show that the optimum system welfare can be approximated by a decentralized algorithm

¹We refer to nodes, devices, and users, interchangeably.

based on Metropolis-Hastings sampling. The algorithm requires the knowledge about the fraction of devices that forward a given channel which is estimated from local observations by individual devices. We also identify a specific class of system welfare objectives for which the Metropolis-Hastings algorithm does not require as input the fraction of devices that forward a given channel, and thus no estimation is needed. We show simulation results that confirm that by using our proposed algorithms the system welfare concentrates near the optimum and that rates of convergence are of interest in practice.

Our contributions can be summarized in the following points:

- We propose a framework for optimizing the dissemination of multiple information channels in wireless ad-hoc networks. The optimization is with respect to dissemination times of individual channels subject to the end-user cache capacity constraints. To the best of our knowledge, this is the first proposal for optimizing dissemination of *multiple* information channels in wireless ad-hoc networks with respect to a global system welfare objective. The framework enables (1) direct engineering by allowing derivation of the algorithms for a given system welfare objective, and (2) reverse engineering in that for a given channel selection algorithm deployed by individual devices, one can determine the underlying system welfare objective.

- We show that an optimum system assignment of channels and users can be found by a centralized greedy algorithm.

- We show that the optimum system welfare can be well approximated by a decentralized algorithm based on Metropolis-Hastings sampling, run by individual devices using only local observations.

- We show how to incorporate the objective of optimizing the battery expenditure.

- We present extensive simulation results that validate and demonstrate practicality of the proposed algorithms.

Outline of the paper. Sec. II introduces the system model and notation. Sec. III presents modeling and empirical analysis that characterize the relation between the dissemination time and the fraction of devices that forward a given channel. In this section, we also define the system welfare objective, the utilities associated to channels, and discuss some of their properties. Sec. IV introduces the system welfare problem and contains the result that the problem can be solved by a centralized greedy algorithm. Sec. V quantifies the optimum system welfare and compares with some previously-proposed heuristics using Zune data. Sec. VI presents our Metropolis-Hastings algorithm. In Sec. VII, we show simulation results. Finally, Sec. VIII discusses related work and Sec. IX concludes the paper. We defer some of the proofs to Appendix [6].

II. SYSTEM MODEL AND NOTATION

We consider a system of N wireless nodes participating in the ad-hoc dissemination of J channels. We denote with $\mathcal{U}$ and $\mathcal{J}$ the sets of users and channels, respectively. Every node u is subscribed to a set $S(u)$ of channels, which is assumed

to be fixed. In addition, every node maintains a set of *helped* channels, i.e. channels whose content this node keeps in its public cache in order to facilitate their dissemination.

At a contact of a pair of nodes, the nodes update the content of their respective caches. More precisely, if a pair of nodes u and u' meet, u receives from u' the content that is newer at u' for the channels that u either subscribes or helps, and vice-versa. We do not account for the overhead of establishing contacts and negotiating content updates. We assume that when nodes meet the contact duration is large enough for all useful content to be exchanged, i.e. we assume that the bottlenecks are disconnection times and availability of content at node caches. In addition, we assume that with some rate, every node fetches the content through direct access to the Internet contact for the subscribed and helped channels.

At a given time instant, we call x the global system configuration, defining the assignment of users to channels, i.e.

$$x_{u,j} = 1 \Leftrightarrow \text{node } u \text{ subscribes or helps channel } j.$$

let $F(u,x)$ be the set of channels forwarded by user u and let $H(u,x)$ be the set of channels helped by node u under system configuration x . Indeed,

$$F(u,x) = H(u,x) \cup S(u), \quad u \in \mathcal{U}.$$

We assume that every node u has a cache that can store content for at most C_u channels where $C_u \geq |S(u)|^2$, thus every node can store all the subscribed channels. The configuration x is thus constrained by

$$|F(u,x)| \leq C_u \text{ for all } u \in \mathcal{U}.$$

The problem is to find a configuration x that satisfies the latter constraint and maximizes an appropriately defined global system objective (defined in the next section). Furthermore, we want to find a decentralized algorithm that would yield nearly optimal system configuration which is done in Sec. VI.

We use the following notation:

$$\begin{aligned} s_j &= \text{the fraction of users subscribed to channel } j \\ f_j(x) &= \text{the fraction of users forwarding channel } j \\ &= \frac{1}{N} \sum_{u \in \mathcal{U}} x_{u,j}. \end{aligned}$$

Without loss of generality and unless indicated otherwise, we assume that channels are labeled in a non-increasing order with respect to the channel subscription popularity, i.e. $s_1 \geq s_2 \geq \dots \geq s_J$. Finally, we will use the following notation $\vec{s} = (s_1, \dots, s_J)$ and $\vec{f} = (f_1, \dots, f_J)$.

III. DISSEMINATION TIME AND UTILITY

In order to understand the relation between the dissemination time of a channel and the fraction of nodes forwarding the channel, we first derive an explicit relationship for random mixing mobility model. We then consider the same by empirical analysis of some real-world mobility traces.

² $|A|$ denotes the cardinality of a finite set A .

A. Model-Based Dissemination Time

We consider an arbitrary channel j and suppose that at the time origin ($t = 0$), a piece of content of channel j is generated. We assume that the system configuration x is fixed and, for simplicity, in this section omit it from the notation. Let $\sigma_j(t)$ be the fraction of nodes subscribed to channel j that at time t have received a piece of content generated by a source of channel j at a time instant in $[0, t]$. Similarly, let $\phi_j(t)$ be the fraction of nodes forwarding channel j that have received a piece of content generated by a source of channel j at a time instant in $[0, t]$. The system dynamics is fully described by the following system of ordinary differential equations:

$$\frac{d}{dt}\sigma_j(t) = (\lambda_j + \eta\phi_j(t))(s_j - \sigma_j(t)) \quad (1)$$

$$\frac{d}{dt}\phi_j(t) = (\lambda_j + \eta\phi_j(t))(f_j - \phi_j(t)) \quad (2)$$

where λ_j is the contact rate between every node and a source of channel j and η is the contact rate between nodes. These equations hold under the “random node mixing” assumption and are exact asymptotic for large N . It is not difficult to find that system (1)-(2) admits a closed-form solution (details can be found in [6]). In particular, we have

$$\sigma_j(t) = \sigma_j(0) + (s_j - \sigma_j(0)) \cdot \frac{(\lambda_j + \eta\phi_j(0))(1 - e^{-(\eta f_j + \lambda_j)t})}{\lambda_j + \eta\phi_j(0) + \eta(f_j - \phi_j(0))e^{-(\eta f_j + \lambda_j)t}}. \quad (3)$$

Given $0 < \alpha < 1$, let t_j be the smallest time at which the fraction of subscribers to channel j receives a piece of content generated by a source of channel j at the time origin 0 or later, i.e. at a time instant in $[0, t_j]$. We refer to t_j as the dissemination time for channel j and this the metric of our primary interest. Using (3) and $\sigma_j(t_j) = \alpha$, we obtain the following closed-form expression for the dissemination time

$$t_j = \frac{1}{\lambda_j + f_j\eta} \ln \frac{(f_j - \phi_j(0))\eta K_j + \lambda_j + \eta\phi_j(0)}{(\lambda_j + \eta\phi_j(0))(1 - K_j)} \quad (4)$$

where

$$K_j = \frac{\alpha - \frac{\sigma_j(0)}{s_j}}{1 - \frac{\sigma_j(0)}{s_j}}.$$

We note the following fact whose proof is available in Appendix [6].

Proposition III.1. *The dissemination time t_j in Eq. (4) is a non-increasing, strictly convex function of f_j .*

A special regime of interest is when the content is fetched by direct access to a source at a small rate, that is when the dissemination time is predominantly determined by the epidemic dissemination. In this case, we have $\sigma_j(0) \ll \lambda_j/\eta \ll s_j$ and $\phi_j(0) \ll \lambda_j/\eta \ll f_j$ and Eq. (4) becomes

$$t_j \approx \frac{1}{\eta f_j} \left(\ln \frac{\alpha}{1 - \alpha} + \ln \frac{\eta}{\lambda_j} + \ln f_j \right). \quad (5)$$

In this case, the dissemination time is approximately inversely

Fig. 1. The dissemination time versus the fraction of forwarding nodes inferred from the CAM data. Each mark shows the median value of the dissemination time obtained by taking each node as a source and repeating for 10 random samples of the forwarding nodes.

proportional to the fraction of nodes forwarding the channel up to a multiplicative factor that is logarithmic in the fraction of nodes that forward the channel.

B. Empirical Dissemination Time

We infer the dissemination time from some real-world mobility traces. Specifically, we consider (1) traces of contacts between human-carried devices in Cambridge area, UK [7] and (2) traces of taxi positions in the area of San Francisco [8]. We refer to these two datasets as CAM and SF-TAXI, respectively. The CAM dataset contains information for 36 Bluetooth-equipped devices over slightly more than 10 days. SF-TAXI contains GPS coordinates for slightly more than 500 taxis over a one month period; we define a pair of taxis to be in contact if the distance between them is smaller or equal to 500 meters [9].

The dissemination time is inferred by the following procedure: (1) we fix a fraction of nodes α picked uniformly at random and designate them as forwarders of the channel; (2) we then inject a piece of content to a randomly selected forwarder and then pass through the trace recording instances at which a forwarder first receives the piece through a contact to a forwarder that already received the piece; (3) the smallest time at which all the forwarders received the piece is declared as the dissemination time. The procedure is repeated for 10 sample sets of forwarding nodes.

In Fig. 1 we show the median dissemination time computed by the above described procedure versus the fraction of forwarding nodes for CAM trace (similar results are obtained for SF-TAXI but are omitted for space reasons; available in Fig. 2 [6]). We observe that the dissemination time is well fitted by a diminishing-decrease curve.

C. The System Welfare

We assume that for every channel j there is an associated utility function $U_j(t_j)$ that quantifies the satisfaction of a subscriber with receiving the content of this channel with the dissemination time t_j . Notice that it is natural to assume that $U_j(t_j)$ is a non-increasing function of t_j .

We denote with $V_j(f_j) = U_j(t_j(f_j))$ the utility function for channel j with respect to the fraction of users forwarding channel j . Notice that it is natural to assume that $V_j(f_j)$ is a non-decreasing function of f_j . This indeed follows if both $U_j(t_j)$ and $t_j(f_j)$ are no-increasing functions.

We consider as system welfare objective, a weighted sum of the utilities of individual channels, i.e. for a given set of non-negative weights $\vec{w} = (w_1, \dots, w_J)$, we define the system welfare function $V(\vec{f})$ as follows

$$V(\vec{f}) = \sum_{j \in \mathcal{J}} w_j V_j(f_j).$$

In particular, there are two special cases of interest, each corresponding to different fairness objectives. Case 1 (*channel centric*) where each weight w_j takes value 1, hence, we have

$$V_{CH}(\vec{f}) = \sum_{j \in \mathcal{J}} V_j(f_j) \quad (6)$$

where V_j is a per-channel metric, for example as in Eq.(4) or Eq.(5). Case 2 (*user centric*) where w_j is set proportional to the fraction of subscribers to channel j , hence,

$$V_{US}(\vec{f}) = \sum_{j \in \mathcal{J}} s_j V_j(f_j). \quad (7)$$

Notice that in Sec. VI, we will show that the utility framework can easily be extended to account for the battery energy.

Concave Utility Functions: In Section IV we consider a system welfare problem that can be solved by a greedy algorithm under assumption that (C) $V_j(f_j)$ is a concave function with respect to f_j for every channel j . We identify a set of *sufficient* conditions for condition (C) to hold true.

Proposition III.2. *Suppose that following two conditions hold*

- (C1) $U_j(t_j)$ is a non-increasing, concave function of t_j ,
- (C2) $t_j(f_j)$ is a convex function of f_j .

Then, $V_j(f_j)$ is a concave function of f_j .

The result follows by elementary convexity properties that are described in [6]. Condition (C1) means that the subscriber's dissatisfaction is increasing with the dissemination time t_j . This conforms to the case where the subscriber would like to receive the fresh content within a time horizon of its generation and becomes increasingly dissatisfied as the delivery time increases beyond the time horizon. Condition (C2) means that the dissemination time $t_j(f_j)$ exhibits diminishing returns with increasing the fraction of nodes f_j forwarding channel j ; recall that Sec. III-A and Sec. III-B support this assumption.

IV. THE SYSTEM WELFARE PROBLEM

We consider the system welfare problem where the objective is to optimize the aggregate utility of channel dissemination times subject to the user capacity constraints. Solving the system welfare problem corresponds to finding an *assignment*

of users to channels that solves the following problem:

SYSTEM

$$\begin{aligned} & \text{maximize} && \sum_{j=1}^J w_j V_j \left(\frac{1}{N} \sum_{u=1}^N x_{u,j} \right) \\ & \text{over} && x_{u,j} \in \{0, 1\} \\ & \text{subject to} && \sum_{j=1}^J x_{u,j} \leq C_u \\ & && x_{u,j} = 1, (u, j) : j \in S(u). \end{aligned}$$

Notice that defining the system welfare as a sum of individual utilities is rather standard, in particular, in the context of resource allocation. The weights $(w_1, w_2, \dots, w_J)$ are non-negative real numbers that can be arbitrarily fixed. In particular, it is of appeal to let w_j be proportional to the fraction of subscribers to channel j as in this case, the utility $V_j(\cdot)$ can be interpreted as the utility function for channel j for a typical subscriber to channel j .

In the following, we establish that we can rephrase SYSTEM as an equivalent optimization over the number of helpers per individual channels. Let H_j be the number of users who help channel j and let $\vec{H} = (H_1, H_2, \dots, H_J)$. Let us define $v(A)$ for $A \subseteq \mathcal{J}$ by

$$v(A) = \sum_{u \in \mathcal{U}} \min \left(\sum_{j \in A} 1_{j \in S(u)}, C_u - |S(u)| \right) \quad (8)$$

and let $P(v)$ be the polyhedron defined by

$$P(v) = \left\{ x \in \mathbb{N}_0^J : \sum_{j \in A} x_j \leq v(A), A \subseteq \mathcal{J} \right\}.$$

We consider the following problem:

SYSTEM-H

$$\begin{aligned} & \text{maximize} && \sum_{j=1}^J w_j V_j \left(s_j + \frac{1}{N} H_j \right) \\ & \text{over} && \vec{H} \in P(v). \end{aligned}$$

In the following theorem, we establish a relation between SYSTEM and SYSTEM-H.

Theorem IV.1. *The optimal value of a solution to SYSTEM is equal to that of SYSTEM-H.*

Proof: Proof is based on a reduction to a max-flow problem and is available in Appendix [6]. ■

We next assert conditions under which SYSTEM-H can be solved by a greedy allocation of channels to users. Let us denote with $\Delta_j V(\vec{s} + \frac{1}{N} \vec{H})$ the increment of the aggregate utility by assigning a user to channel j , i.e.

$$\begin{aligned} \Delta_j V \left(\vec{s} + \frac{\vec{H}}{N} \right) &= V \left(\vec{s} + \frac{\vec{H} + e_j}{N} \right) - V \left(\vec{s} + \frac{\vec{H}}{N} \right) \\ &= w_j \left[V_j \left(s_j + \frac{H_j + 1}{N} \right) - V_j \left(s_j + \frac{H_j}{N} \right) \right] \end{aligned}$$

where e_j is a vector of dimension $|\mathcal{J}|$ with all the coordinates

equal to 0 but the j th coordinate equal to 1. We consider the following greedy assignment.

Algorithm 1 Centralized GREEDY Algorithm for Allocation of Helped Channels.

```

1:  $H = 0$ 
2: while 1 do
3:   Find  $I \in \mathcal{J}$  such that  $\vec{H} + e_I \in P(v)$ 
4:   and  $\Delta_I V(\vec{s} + \vec{H}/N) \geq \Delta_j V(\vec{s} + \vec{H}/N)$  for all  $j \in \mathcal{J}$ 
5:   such that  $\vec{H} + e_j \in P(v)$ 
6:
7:   if there exists no such  $I$  then break
8:   end if
9:    $H_I \leftarrow H_I + 1$ 
10: end while

```

Theorem IV.2. Suppose that $V_j(x)$ is a concave function for every $j \in \mathcal{J}$, then GREEDY finds a solution to SYSTEM.

Proof is available in [6]. The result follows from Theorem IV.1 and showing that SYSTEM-H is a maximization of a concave function over a submodular polyhedron. This ensures that conditions for optimality of the greedy procedure in [10] are met.

A. Some Particular Channel Selection Strategies

We introduce three particular channel selection strategies. The first two strategies are closely related to the uniform and most solicited strategies in [1]. The third strategy arises from the Metropolis sampling in Sec. VI.

Uniform: Each user u picks a subset of $C_u - |S(u)|$ channels by sampling uniformly at random without replacement from the set of channels to which user u is not subscribed, i.e. from the set $\mathcal{J} \setminus S(u)$.

Notice that this strategy biases to forwarding less popular channels. This is indeed intuitive as each user selects from the subset of channels that contains channels not subscribed by this user, so channels that have fewer subscribers would have larger chance of being selected. The interested reader is referred to [6] for more details.

Top Popular: Each user u picks a set of most subscribed channels from the set of channels $\mathcal{J} \setminus S(u)$. This scheme could be seen as a greedy scheme for biasing to popular channels; we will consider this scheme in numerical evaluations in Sec. V.

A Pick from a Neighbor: This is a class of channel selection strategies at an encounter of user u and u' , user u picks a candidate channel from user those at user u' , and then based on some decision process decides whether to replace a currently helped channel with the candidate channel. The decision process is assumed to be local and independent of current assignment of users to channels, which makes these strategies of practical interest.

In Sec. VI, we will construct one such a scheme and show that it is associated with to the system welfare problem with the following objective function: $V_{PFN}(\vec{f}) = \sum_{j \in \mathcal{J}} V_j^{PFN}(f_j)$

Fig. 2. The dissemination time versus the helper cache size C . (For user u , the total cache size is C_u so that $C_u = |S(u)| + C$ where $|S(u)|$ is the number of subscribed channels by user u .)

with

$$V_j^{PFN}(f_j) = (a_j + C)f_j + Df_j \ln f_j \quad (9)$$

where C and D are global system constants and $a_j \geq 0$ is a constant for channel j , which signifies the relative importance of this channel. Note that function $V_j^{PFN}(f_j)$ in Eq. (9) is non-decreasing with f_j . Note, however, that $V_j^{PFN}(f_j)$ is a convex function of f_j . Hence, it does not validate the conditions discussed in Sec. III-C that ensure optimality of the greedy assignment in Sec. IV.

V. SYSTEM OPTIMUM VS. HEURISTICS

In this section we provide a support for the following claim: The optimum system welfare can be significantly larger than that of heuristic assignments of users to channels suggested in prior work. In particular, we compare with the Uniform and Top Popular assignments defined in the preceding section.

We define user subscription sets according to the subscriptions to podcast feeds observed in the Zune Social dataset. This dataset provides subscription information for more than a million of users over slightly more than 8,000 podcasts feeds. The distribution of the number of subscriptions per user is quite skewed with the median value of 3 and the mean value of about twice that value [5].

We consider the user-centric system welfare with the channel utility functions $V_j(f_j) = -t_j(f_j)$ where $t_j(f_j)$ is the dissemination time given by Eq. (4). For every user u , we set $C_u = |S(u)| + C$ where $|S(u)|$ is specified by the input data and $C \geq 0$ is a parameter. We compute an optimum assignment by the algorithm GREEDY (Sec. IV). Uniform and Top Popular assignments are computed as prescribed by their respective definitions.

In Fig. 2, we show the dissemination time versus the helper cache size C . For every user, the rate of the access to the infrastructure is fixed to 1 access per day. The rate at which every user encounters other users is fixed to 100 per day. If the dissemination is solely by a direct access to the infrastructure, then the dissemination time is slightly more than half a day (about 13.5 hours). In Fig. 2, we observe that the dissemination time under a system-welfare optimum assignment can be reduced for as much as almost 50% using the device-to-device

Fig. 3. The dissemination time versus the rate of encounters η . The cache size of user u is set to $C_u = |S(u)| + C$ with $C = 20$.

dissemination. Perhaps even more interestingly, we observe that the gap between the system-welfare optimum and that of the Uniform and Top Popular assignments can be significant.

Furthermore, in Fig. 3 we present the results under the same setting as in Fig. 2 but varying the encounter rate and keeping the cache size C fixed to 20. These results show a lack of order for the Uniform and Top Popular assignments; for some cases one is better than the other one, and vice-versa in other cases. In any case, system-welfare optimum indeed provides best performance.

VI. A DISTRIBUTED METROPOLIS HASTINGS ALGORITHM

In this section we consider the problem of designing a decentralized algorithm. The goal for every node is to decide which channels to help disseminating so that the resulting global configuration x is near optimum system welfare given by

$$V(x) = \sum_{j \in J} w_j V_j(f_j(x)). \quad (10)$$

Notice that unlike to Sec. III we make the dependence on the global configuration x explicit.

A. Metropolis-Hastings

We propose to use the Metropolis-Hastings sampling [11] as it lends itself to distributed optimization, and was successfully used in distributed control problems in wireless networks (e.g. [12]). Before describing our distributed algorithm, we first give a short description of a centralized version of the Metropolis-Hastings algorithm:

At every time step, the algorithm picks a tentative configuration x' , with probability $Q(x, x')$, where x is the current configuration. We assume that the matrix $Q(x, x')$ has the weak symmetry property:

$$Q(x, x') > 0 \Rightarrow Q(x', x) > 0$$

for all $x \neq x'$. The tentative configuration is accepted (i.e. becomes the new configuration) with probability $p = \min(1, q)$ where

$$q = \frac{\pi(x')Q(x', x)}{\pi(x)Q(x, x')} \quad (11)$$

and $\pi(\cdot)$ is a probability distribution on the set of feasible configurations. The algorithm does not converge in the strict sense, however, as the number of iterations tends to be large, the probability distribution of the configuration x converges to the a priori distribution $\pi(\cdot)$. Typically, one uses for $\pi(\cdot)$ a Gibbs distribution given by

$$\pi(x) = \frac{1}{Z} e^{\frac{V(x)}{T}} \quad (12)$$

where T is a system parameter (commonly referred to as a “temperature”) and Z is the normalizing constant. If T is small, the distribution $\pi(\cdot)$ is very much concentrated on the large values of $V(x)$, so that the algorithm produces random configurations that tend to maximize $V(x)$.

B. A Distributed Rewiring Algorithm

We use the Metropolis-Hastings sampling as follows. We use a Gibbs distribution, as in Eq. (12) with $V(\cdot)$ the system-welfare function in Eq. (10). We consider every contact between two nodes as one step of the algorithm. When two nodes meet, they opportunistically exchange content updates; then one of them, say u is selected as a leader and attempts to replace one of its helped channels by one of the channels forwarded from the set held by the other node, say v , as described in Algorithm 2. We now turn to the computation of

Algorithm 2 Distributed Algorithm for Allocation of Helped Channels

- 1: **if** $F(v, x) \subset F(u, x)$ **then** do nothing
- 2: **else**
- 3: u selects one channel j uniformly at random in the set $H(u, x)$
- 4: u selects one channel j' uniformly at random in the set $F(v, x) \setminus F(u, x)$
- 5: compute the acceptance probability $p = \min(1, q)$ with q given by Eq.(15)
- 6: draw a random number U uniformly in $[0, 1]$;
- 7: **if** $U < p$ **then** drop channel j and adopt channel j' as a helped channel
- 8: **end if**
- 9: **end if**

the acceptance probability (line 5 of the algorithm), as given by Eq.(11). First, we compute $Q(x, x')$ where $x' = x - e_{u,j} + e_{u,j'}$ is the new configuration ($e_{u,j}$ is a configuration with all elements equal to 0 but the (u, j) -th element equal to 1).

Proposition VI.1. *The following holds*

$$\frac{Q(x', x)}{Q(x, x')} = \frac{\sum_{v \neq u} \frac{1_{j \in F(v, x)}}{|F(v, x) \setminus F(u, x)| + 1_{j' \notin F(v, x)}}}{\sum_{v \neq u} \frac{1_{j' \in F(v, x)}}{|F(v, x) \setminus F(u, x)|}}. \quad (13)$$

Proof is provided in Appendix [6]. Notice the following approximation

$$\frac{Q(x', x)}{Q(x, x')} \approx \frac{f_j(x)}{f_{j'}(x)}. \quad (14)$$

We next note the following fact:

Proposition VI.2. *Suppose that for a constant $D > 0$, $\lim_{N \rightarrow +\infty} NT = D$. Then,*

$$\lim_{N \rightarrow +\infty} \frac{V(x') - V(x)}{T} = \frac{1}{D} \left(w_{j'} V_{j'}'(f_{j'}(x)) - w_j V_j'(f_j(x)) \right).$$

Proof is available in Appendix [6]. In view of the last proposition, we have

$$\begin{aligned} q &= \frac{Q(x', x)}{Q(x, x')} e^{\frac{1}{T}(V(x') - V(x))} \\ &\approx \frac{Q(x', x)}{Q(x, x')} e^{\frac{1}{NT} (w_{j'} V_{j'}'(f_{j'}(x)) - w_j V_j'(f_j(x)))}. \end{aligned}$$

Combining with (14) we obtain for q the value

$$q = \frac{f_j(x)}{f_{j'}(x)} e^{\frac{1}{D} (w_{j'} V_{j'}'(f_{j'}(x)) - w_j V_j'(f_j(x)))} \quad (15)$$

where $D = NT$ is a global system parameter.

Algorithm 2 requires node u to estimate f_j and $f_{j'}$. This can be done by having the nodes exchange, when they meet, updates of channel popularity for all channels that they know of and then performing an exponential smoothing. A simple, but memory demanding scheme is as follows. Every node u maintains for every channel j an estimate $\hat{f}_j$. When node u meets node u' , for all channels that u' helps or subscribes to, node u does $\hat{f}_j \leftarrow a + (1-a)\hat{f}_j$ and for all other channels $\hat{f}_j \leftarrow (1-a)\hat{f}_j$ where $0 < a < 1$.

Furthermore, all nodes need to share the global system variable D and know the utility function of each channel where the latter can be contained as a meta-information in the channel data. In the following section, we give an algorithm for a particular choice of the system-welfare function, which does not require such estimations.

C. A Simplified Algorithm

It is possible to entirely alleviate the estimation of the f_j variables, albeit at the expense of imposing a family of utility functions. The idea is to choose a set of utility functions $V_j(\cdot)$ such that f_j and $f_{j'}$ cancel out in Eq. (15). This results in a scheme that belongs to the class of schemes "A pick from a neighbor" that we discussed in Sec. IV.

Theorem VI.1. *Suppose that for every channel j , the utility function is $V_j^{PFN}(\cdot)$ in Eq. (9), then q in Eq. (15) is given by*

$$q = \frac{\beta_{j'}}{\beta_j} \quad (16)$$

where $\beta_j = e^{\frac{a_j}{D}}$ and $\beta_{j'} = e^{\frac{a_{j'}}{D}}$. In particular, q is independent of $f_j(x)$ and $f_{j'}(x)$, and more generally of the configuration x .

Proof: Follows from Eq. (9) and Eq. (15). ■

With this simplified algorithm, nodes need to know the static parameters $\beta_j > 0$ associated with each channel. There is no global constant, nor it is necessary to evaluate $f_j(x)$. Higher values of j mean that we give more value to disseminating channel j more quickly. Note that only the relative values of

β_j matter, as Eq. (16) uses only the ratios, and β_j can thus be interpreted as the priority level for channel j . The resulting algorithm is the same as Algorithm 2 but with the acceptance probability q computed using Eq. (16) instead of Eq. (15).

D. A Battery Saving Algorithm

The previous algorithm may be improved to account for battery saving. The motivation is that a node may be reluctant to help disseminate channels if its battery level is low. We address this issue as follows. Assume that every node u knows its battery level $b_u \geq 0$. The battery is empty when $b_u = 0$. Assume to simplify that all nodes measure b_u in the same scale, for example, number of remaining hours of operation at full activity. We can replace the global utility in Eq. (10) by

$$\sum_{j \in \mathcal{J}} w_j V_j(f_j) - \sum_{u \in \mathcal{U}} W_u(b_u)$$

where $W_u(b)$ is a convex, decreasing function of b (for example $W_u(b) = \frac{1}{b^m}$), such that $W_u(b)$ expresses the penalty perceived by user u when its battery level is b . We can apply the Metropolis-Hastings algorithm with this new function. The only difference is in the computation of the acceptance probability. The computation of q in Eq. (16) is replaced by

$$q = \frac{\beta_{j'}}{\beta_j} e^{-[h_u(b_u) - h_{u'}(b_{u'})]} \quad (17)$$

where u and u' are the two nodes involved in the interaction and $h_u(b) > 0$ is the marginal cost of exchanging a channel when two nodes meet, divided by the temperature T (an increasing function of b).

The resulting algorithm is the same as Algorithm 2 with Eq. (15) in line 5 replaced by Eq. (17). The required configuration is (1) every channel j has a static priority level $\beta_j > 0$ and (2) every node u knows its own function $h_u(b)$ for the cost of exchanging one channel with a neighbor when this node's battery level is b .

VII. SIMULATION RESULTS

In this section we present simulation results whose purpose is: (1) to demonstrate that the system welfare under the distributed Metropolis-Hastings algorithm is near to the optimum system welfare, and (2) to demonstrate cases where optimizing the system welfare under a real-world mobility results in good system configurations.

In order to cover a broad set of parameters, we conducted simulations by varying the parameters along the following dimensions: (1) the node mobility is either random mixing or using a real mobility trace, (2) the system size is either small or large with respect to the number of users and the number of channels, (3) different distributions of subscriptions per channel are used, (4) the fractions of nodes forwarding a given channel is either globally known or locally estimated, and (5) a range of temperature parameters for the Metropolis-Hastings system is used. We consider the random mixing mobility in order to provide results for scenarios for which we well understand the relation between the channel dissemination

Fig. 4. Convergence of the Metropolis-Hastings algorithm under channel-centric system welfare: (a) small-scale, Zipf-2/3, (b) small-scale, Zipf-1, (c) large-scale, Zipf-2/3, (d) large-scale, Zipf-1. Small-scale refers to $(N, J) = (20, 20)$ and the large-scale refers to $(N, J) = (200, 100)$. The y-axis is the mean dissemination time over all channels. The thick horizontal line denotes the system optimum mean dissemination time. Other solid curves are for the Metropolis-Hastings algorithm with the portions of nodes that forward any given channel known ($\vec{f}$). The dashed lines denote the same but with $\vec{f}$ locally estimated.

time and the fraction of forwarding nodes of this channel. We used our own discrete-event simulator.

A. Random-Mixing Mobility

We assume a random mixing mobility where the dissemination time of a channel is a function of the fraction of nodes forwarding the channel with the exact relation established in Sec. III-A.

We consider a small size system where the number of user is equal to 20 and the number of channels is also equal to 20. We also consider a large size system where the number of users is 200 and the number of channels is equal to 100. For the fractions of subscribers per channel $\vec{s}$, we assume a Zipf distribution with the scale parameter either 2/3 or 1. We consider the Zipf distribution for the following reasons: (1) it well captures subscription patterns that typically exhibit a power law tail, (2) it was considered in previous work [1], and (3) it allows for reproducibility of results. For the system welfare objective, we consider channel-centric and user-centric cases with the utility function $V_j(f_j) = -t_j(f_j)$ for channel j where $t_j(f_j)$ is the dissemination time and f_j is the fraction of nodes forwarding channel j . For the function $t_j(f_j)$ we admit Eq. (4). In cases when $\vec{f}$ or $\vec{s}$ are locally estimated, each node uses an exponential weighted averaging with the smoothing constant (1) equal to 0.9 for the estimation of $\vec{f}$ and (2) equal to 0.1 and 0.02 for the estimation of $\vec{s}$ under channel-centric and user-centric case, respectively.

In Fig. 4 we show the results obtained for the channel-centric case. The graphs show the mean dissemination time per channel, i.e. $(\sum_{j \in \mathcal{J}} t_j(f_j))/J$, versus the number of encounters per node. We show the results for the Metropolis-Hastings with $\vec{f}$ assumed to be either known or locally estimated by individual nodes. We observe that the system welfare under the Metropolis-Hastings algorithm concentrates near the optimum system welfare. The results in Fig. 4 indicate a faster concentration in cases when $\vec{f}$ is globally known. We obtained qualitatively same results for the user-centric case which we omit for space reasons; the interested reader is referred to Fig. 9 in [6]. In summary the results support that the system welfare under the Metropolis-Hastings algorithm concentrates

TABLE I
THE DISSEMINATION TIME PER CHANNEL AND PER USER IN MINUTES FOR CAM TRACE.

Channel-centric	UNI	TOP	OPT
Median	70.2500	133.1000	52.1429
Mean	70.4700	137.1250	57.2000
User-centric	UNI	TOP	OPT
Median	70.4028	97.4528	56.9333
Mean	70.0578	102.7284	59.4089

near the optimum system welfare with $\vec{f}$ (and $\vec{s}$ in the user-centric case) either globally known or locally estimated.

B. Real-Trace Mobility

We compare the system performance under an optimum assignment of channels to users (OPT) with that of the heuristics Uniform (UNI) and Top Popular (TOP) which were introduced in Sec. IV-A. We define the system welfare objective using the dissemination function $t_j(f_j)$ inferred from the mobility trace CAM and assuming $V_j(f_j) = -t_j(f_j)$ as in the preceding section. Specifically, we define the logarithm of $t_j(f_j)$ by a concatenation of linear segments that closely follow the empirical data (available in Fig. 7 [6]). We consider a scenario with $J = 40$ channels, 10 subscriptions per each user, and 10 channels helped by each user. We assume that the channel subscription rates follow a Zipf distribution with the scale parameter 2/3. For each setting of the simulation parameters, we repeat the experiment 5 times, each time injecting a message of a channel to a user picked uniformly at random from the users forwarding the channel at the beginning of the trace. Notice that CAM trace is for 36 distinct users with the encounter rate η equal to 0.001 per second, i.e. 1.2 users every two minutes.

In Table 1 we show the median and mean dissemination time per channel and per user for the channel-centric and user-centric cases, respectively. For both the mean and the median dissemination time, OPT substantially outperforms UNI and TOP for in channel-centric and user-centric case. In particular, in the channel-centric case, OPT achieves smaller dissemination for as much as 70 minutes than TOP and for

Fig. 5. The mean channel dissemination time under CAM mobility trace with the channel-centric system welfare. Channels are enumerated in decreasing order of subscription popularity.

as much as 10 minutes than UNI for both the mean and the median dissemination time. In the user-centric case, OPT achieves the dissemination time that is smaller for as much as 40 minutes than TOP and for as much as 10 minutes than UNI for both the mean and the median dissemination time. Furthermore, in Fig. 5, we show the mean dissemination time for each channel. We observe the following (1) under the channel assignment UNI, some intermediate popular channels may be penalized with a high dissemination time; in particular, in Fig. 5, we note that the tenth most popular channel obtains as much as five hours larger dissemination time than under other channel assignments; and (2) the same can happen under TOP where the results conform to the expected bias against less popular channels. Notice that the dissemination time for several less popular channels is larger or several hours than under other channel assignments. In summary, the results demonstrate cases where assigning channels by optimizing the system welfare avoids penalizing some channels which can occur under heuristic assignments such as UNI or TOP.

VIII. RELATED WORK

In this section we discuss the work that is most closely related to our work; more discussion is available in [6]. [1] proposed several heuristics for the content exchange between devices based on the inferred preference of the user of a device and that of encountered devices. Each device is assumed to forward an unlimited number of feeds and prioritizes the download of pieces of the content feeds from the encountered devices. Feeds subscribed by a device take priority over other feeds. Each device uses a solicitation strategy to decide which pieces to fetch from the encountered devices. The considered solicitation strategies include the *most solicited* and *uniform* which essentially correspond to the top popular and uniform channel assignment considered in our paper. The approach in [1] is different from ours in that they evaluate a set of a priori defined strategies while we first formulate a global system welfare objective and then identify a channel selection strategy to optimize the given system objective. Another closely related

work is CarTorrent [2], a peer-to-peer file sharing designed for vehicular network scenarios using epidemic-style content dissemination. Further related work is [13]. Our work is generally different from the state-of-the-art results on epidemic-style information dissemination in that our goal is efficient dissemination by controlling the epidemic spread of *multiple* content streams.

IX. CONCLUSION

We proposed a framework for optimizing dissemination of multiple channels in wireless networks that leverages both access direct access to the Internet and information transfer between mobile devices. The problem amounts to finding an optimum assignment of channels to users for forwarding the content with respect to a given global system welfare objective. We showed that an optimum assignment can be found by a centralized greedy algorithm. Moreover, we showed that near optimum system welfare can be achieved by a decentralized algorithm based on Metropolis-Hasting sampling. We also showed how to account for the battery energy.

There are several interesting directions for future work like (1) conducting a more thorough analysis of the convergence of the proposed two-timescale control (one timescale for the Metropolis-Hastings rewiring of the system configuration and other timescale for the estimation of the fractions of forwarding nodes per channel) and (2) extending the framework to mobility scenarios where the encounter rates between mobile nodes can be heterogeneous.

REFERENCES

- [1] V. Lenders, G. Karlsson, and M. May. Wireless Ad Hoc Podcasting. In *Proc. of IEEE SECON*, 2007.
- [2] K. C. Lee, S. H. Lee, R. R. Cheung, U. Lee, and M. Gerla. First Experience with CarTorrent in a Real Vehicular Ad Hoc Network Testbed. *Mobile Networking for Vehicular Env.*, pages 109–114, 2007.
- [3] MIT Reality Mining. Social serendipity. Downloaded from <http://reality.media.mit.edu/serendipity.php>, July 2009.
- [4] Microsoft. Zune social. www.zune.net, December 2008.
- [5] D. Gunawardena, T. Karagiannis, A. Proutiere, and M. Vojnović. Characterizing podcast services: Publishing, usage, and dissemination. In *Proc. of ACM IMC*, 2009.
- [6] L. Hu, J.-Y. Le Boudec, and M. Vojnović. Optimal channel choice for collaborative ad-hoc dissemination. Technical Report MSR-TR-2009-26, Microsoft Research, July 2009.
- [7] J. Scott, R. Gass, J. Crowcroft, P. Hui, C. Diot, and A. Chaintreau. CRAWDAD data set cambridge/haggle (v. 2006-09-15). Downloaded from <http://crawdad.cs.dartmouth.edu/cambridge/haggle>, September 2006.
- [8] M. Piorkowski, N. Sarafijanovic-Djukic, and M. Grossglauser. CRAWDAD data set epfl/mobility (v. 2009-02-24). Downloaded from <http://crawdad.cs.dartmouth.edu/epfl/mobility>, February 2009.
- [9] T. Karagiannis, J.-Y. Le Boudec, and M. Vojnović. Power-Law and Exponential-Decay of Inter-Contact Time Between Mobile Devices. In *Proc. of ACM MOBICOM*, 2007.
- [10] A. Federgruen and H. Groenevelt. The greedy procedure for resource allocation problems: Necessary and sufficient conditions for optimality. *Operations Research*, 34(6):909–918, 1986.
- [11] P. Brémaud. *Markov Chains: Gibbs Fields, Monte Carlo Simulation and Queues*. Springer, 1999.
- [12] B. Kauffmann, F. Baccelli, A. Chaintreau, V. Mhatre, K. Papagiannaki, and C. Diot. Measurement-Based Self Organization of Interfering 802.11 Wireless Access Networks. In *Proc. of IEEE INFOCOM*, 2007.
- [13] A. Chaintreau, J.-Y. Le Boudec, and N. Ristanović. The age of gossip: Spatial mean-field regime. In *Proc. of ACM Sigmetrics*, 2009.